WHAT IS CLAIMED IS:

1. A method comprising:

receiving at least two frames of image data;

learning a model for the appearance of an object from the at least two frames of image data; and

tracking a changing position of the object in three dimensions from the at least two frames of image data.

2. The method of claim 1 wherein each frame of image data consists of image data from at least two cameras.

3. The method of claim 1 wherein tracking a changing position comprises representing possible positions as particles.

4. The method of claim 3 wherein tracking a changing position further comprises weighting each particle in a set of particles based on the probability that the particle represents the position of the object.

5. The method of claim 4 further comprising using the weighting of the particles to resample the particles in the set of particles as part of selecting a set of particles for a next frame.

6. The method of claim 5 wherein selecting a set of particles for a next frame further comprises shifting the resampled particles.

7. The method of claim 5 wherein resampling the particles comprises for each particle of a current frame multiplying the weight of the particle by a total number of desired particles for the next frame to identify a number of resampled particles, and placing that number of resampled particles at the position of the particle for the current frame.

8. The method of claim 1 wherein learning a model for the appearance of an object comprises using an expectation-maximization algorithm to learn the model of the appearance.

9. The method of claim 8 wherein the expectation-maximization algorithm further comprises determining a posterior probability for the appearance of the object.

10. The method of claim 9 wherein determining the posterior probability for the appearance of the object comprises determining the posterior probability based on a distribution having a mean that is a function of a prior model of the appearance of the object, a value in the image data determined from a first camera and a value in the image data determined from a second camera.

11. The method of claim 10 wherein the posterior probability comprises a mixture of distributions.

12. The method of claim 11 wherein one of the distributions in the mixture of distributions is weighted based on the correspondence between a value in the image data determined from a first camera and a value in the image data determined from a second camera.

13. The method of claim 11 wherein one of the distributions in the mixture of distributions is weighted based on the correspondence between at least one value in the image data and a prior model of the appearance of the object.

14. The method of claim 1 further comprising determining a model of the appearance of a background.

15. The method of claim 1 wherein:

tracking a changing position of an object comprises designating each of a set of possible positions for the objects as particles and weighting each particle; and

learning a model for the appearance of an object comprises forming a posterior

probability of the appearance of the object for each particle, weighting each posterior probability based on the weight of the respective particle, and summing the weighted posterior probabilities for the set of particles.

16. A computer-readable medium having computer-executable instructions for performing steps to track the position of an object and to learn a model of the appearance of the object, the steps comprising:

representing possible positions of the object as particles and weighting each particle; and

determining a posterior probability for the appearance of the object based in part on the weighting of the particles.

17. The computer-readable medium of claim 16 wherein determining a posterior probability comprises using image data from at least two cameras.

18. The computer-readable medium of claim 16 wherein determining a posterior probability comprises determining a separate posterior probability for each particle, weighting each separate posterior probability based on the weight of the associated particle, and summing the weighted posterior probabilities to form the posterior probability.

19. The computer-readable medium of claim 18 wherein determining a posterior probability for each particle comprises utilizing a distribution having a mean that is based on a prior model of the appearance of the object and image data from at least one camera.

20. The computer-readable medium of claim 19 wherein determining a posterior probability for each particle further comprises weighting the distribution.

21. The computer-readable medium of claim 20 wherein weighting the distribution comprises applying a weight that is based on the correspondence between an image value from a first camera and an image value from a second camera.

22. The computer-readable medium of claim 20 wherein weighting the distribution comprises applying a weight that is based on the correspondence between a prior model of the appearance and an image value from at least one camera.

23. The computer-readable medium of claim 16 wherein determining a posterior probability forms part of an expectation-maximization algorithm.

24. The computer-readable medium of claim 23 wherein the expectation-maximization algorithm further comprises updating model parameters for a model of the appearance of the object based on the posterior probability.

25. The computer-readable medium of claim 16 further comprising determining a first set of particles for a first frame of image data and determining a second set of particles for a second frame of image data.

26. The computer-readable medium of claim 25 wherein determining a second set of particles comprises selecting the second set of particles based on the weights of the particles in the first set of particles.

27. A method comprising:

receiving image data from a first camera;

receiving image data from a second camera;

selecting a set of particles for a frame of the image data, each particle representing a possible position for an object;

using the image data from the first camera and the second camera to determine a weight for each particle; and

selecting a set of particles for a next frame of the image data based on the weights of the particles.

28. The method of claim 27 further comprising using the image data from the first camera and the second camera to determine a posterior probability for the appearance of a pixel in an object.

29. The method of claim 28 further comprising using the posterior probability to update model parameters that describe the appearance of the pixel.

30. The method of claim 28 wherein determining a posterior probability further comprises determining a separate posterior probability for each particle in the set of particles, weighting each separate posterior probability based on the weights of the respective particles, and summing the weighted posterior probabilities to form the posterior probability.

31. The method of claim 30 wherein determining a posterior probability for a particle comprises determining a mixture of probability distributions.

32. The method of claim 30 wherein determining a mixture of probability distributions comprises determining a distribution having a mean based on a mean for a prior model of the appearance of the

pixel, a value from the image data of the first camera and a value from the image data of the second camera.

33. The method of claim 30 wherein determining a mixture of probability distributions comprises determining a weight for a distribution based on the similarity between a value from the image data of the first camera and a value from the image data of the second camera.

34. The method of claim 30 wherein determining a mixture of probability distributions comprises determining a weight for a distribution based on the degree to which a value in the image data of the first camera and a value in the image data of the second camera match a mean of a prior model for the appearance of the pixel.